


A Classification of Symbolic Transition Systems

Thomas A. Henzinger

Rupak Majumdar



Report No. UCB/CSD-99-1086

December 1999

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE DEC 1999	2. REPORT TYPE	3. DATES COVERED 00-00-1999 to 00-00-1999
4. TITLE AND SUBTITLE A Classification of Symbolic Transition Systems		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER	
	5e. TASK NUMBER	
	5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California at Berkeley, Department of Electrical Engineering and Computer Sciences, Berkeley, CA, 94720		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT <p>We define five increasingly comprehensive classes of infinite-state systems, called STS1-5, whose state spaces have finitary structure. For four of these classes, we provide examples from hybrid systems. STS1: These are the systems with finite bisimilarity quotients. They can be analyzed symbolically by (1) iterating the predecessor and boolean operations starting from a finite set of observable state sets, and (2) terminating when no new state sets are generated. This enables model checking of the mu-calculus. STS2: These are the systems with finite similarity quotients. They can be analyzed symbolically by iterating the predecessor and positive boolean operations. This enables model checking of the existential and universal fragments of the mu-calculus. STS3: These are the systems with finite trace-equivalence quotients. They can be analyzed symbolically by iterating the predecessor operation and a restricted form of positive boolean operations (intersection is restricted to intersection with observables). This enables model checking of linear temporal logic. STS4: These are the systems with finite distance-equivalence quotients (two states are equivalent if for every distance d, the same observables can be reached in d transitions). The systems in this class can be analyzed symbolically by iterating the predecessor operation and terminating when no new state sets are generated. This enables model checking of the existential conjunction-free and universal disjunction-free fragments of the mu-calculus. STS5: These are the systems with finite bounded-reachability quotients (two states are equivalent if for every distance d, the same observables can be reached in d or fewer transitions). The systems in this class can be analyzed symbolically by iterating the predecessor operation and terminating when no new states are encountered. This enables model checking of reachability properties.</p>		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 23	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

A Classification of Symbolic Transition Systems^{*,**}

Thomas A. Henzinger Rupak Majumdar

Department of Electrical Engineering and Computer Sciences
University of California at Berkeley, CA 94720-1770, USA
`{tah, rupak}@eecs.berkeley.edu`

Abstract. We define five increasingly comprehensive classes of infinite-state systems, called STS1–5, whose state spaces have finitary structure. For four of these classes, we provide examples from hybrid systems.

STS1 These are the systems with finite *bisimilarity* quotients. They can be analyzed symbolically by (1) iterating the predecessor and boolean operations starting from a finite set of observable state sets, and (2) terminating when no new state sets are generated. This enables model checking of the μ -calculus.

STS2 These are the systems with finite *similarity* quotients. They can be analyzed symbolically by iterating the predecessor and positive boolean operations. This enables model checking of the existential and universal fragments of the μ -calculus.

STS3 These are the systems with finite *trace-equivalence* quotients. They can be analyzed symbolically by iterating the predecessor operation and a restricted form of positive boolean operations (intersection is restricted to intersection with observables). This enables model checking of linear temporal logic.

STS4 These are the systems with finite *distance-equivalence* quotients (two states are equivalent if for every distance d , the same observables can be reached in d transitions). The systems in this class can be analyzed symbolically by iterating the predecessor operation and terminating when no new state sets are generated. This enables model checking of the existential conjunction-free and universal disjunction-free fragments of the μ -calculus.

STS5 These are the systems with finite *bounded-reachability* quotients (two states are equivalent if for every distance d , the same observables can be reached in d or fewer transitions). The systems in this class can be

* An abbreviated version of this paper will appear in the *Proceedings of the 17th International Symposium on Theoretical Aspects of Computer Science (STACS)*, Lecture Notes in Computer Science, Springer-Verlag, 2000.

** This research was supported in part by the DARPA (NASA) grant NAG2-1214, the DARPA (Wright-Patterson AFB) grant F33615-C-98-3614, the MARCO grant 98-DT-660, the ARO MURI grant DAAH-04-96-1-0341, and the NSF CAREER award CCR-9501708.

analyzed symbolically by iterating the predecessor operation and terminating when no new states are encountered. This enables model checking of reachability properties.

0 Introduction

To explore the state space of an infinite-state transition system, it is often convenient to compute on a data type called “region,” whose members represent (possibly infinite) sets of states. Regions might be implemented, for example, as constraints on the integers or reals. We say that a transition system is “symbolic” if it comes equipped with an algebra of regions which permits the effective computation of certain operations on regions. For model checking, we are particularly interested in boolean operations on regions as well as the predecessor operation, which, given a target region, computes the region of all states with successors in the target region. While a region algebra supports individual operations on regions, the iteration of these operations may generate an infinite number of distinct regions. In this paper, we study restricted classes of symbolic transition systems for which certain forms of iteration, if terminated after a finite number of operations, still yield sufficient information for checking interesting, unbounded temporal properties of the system.

0.1 Symbolic Transition Systems

Definition: Symbolic transition system A *symbolic transition system* $\mathcal{S} = (Q, \delta, R, \ulcorner \cdot \urcorner, P)$ consists of a (possibly infinite) set Q of *states*, a (possibly non-deterministic) *transition* function $\delta : Q \rightarrow 2^Q$ which maps each state to a set of successor states, a (possibly infinite) set R of *regions*, an *extension* function $\ulcorner \cdot \urcorner : R \rightarrow 2^Q$ which maps each region to a set of contained states, and a finite set $P \subseteq R$ of *observables*, such that the following six conditions are satisfied:

1. The set P of observables covers the state space Q ; that is, $\bigcup \{\ulcorner p \urcorner \mid p \in P\} = Q$.
2. For each region $\sigma \in R$, there is a region $Pre(\sigma) \in R$ such that

$$\ulcorner Pre(\sigma) \urcorner = \{s \in Q \mid (\exists t \in \delta(s) : t \in \sigma)\};$$

furthermore, the function $Pre : R \rightarrow R$ is computable.

3. For each pair $\sigma, \tau \in R$ of regions, there is a region $And(\sigma, \tau) \in R$ such that $\ulcorner And(\sigma, \tau) \urcorner = \ulcorner \sigma \urcorner \cap \ulcorner \tau \urcorner$; furthermore, the function $And : R \times R \rightarrow R$ is computable.
4. For each pair $\sigma, \tau \in R$ of regions, there is a region $Diff(\sigma, \tau) \in R$ such that $\ulcorner Diff(\sigma, \tau) \urcorner = \ulcorner \sigma \urcorner \setminus \ulcorner \tau \urcorner$; furthermore, the function $Diff : R \times R \rightarrow R$ is computable.
5. All emptiness questions about regions can be decided; that is, there is a computable function $Empty : R \rightarrow \mathbb{B}$ such that $Empty(\sigma)$ iff $\ulcorner \sigma \urcorner = \emptyset$.

6. All membership questions about regions can be decided; that is, there is a computable function $Member : Q \times R \rightarrow \mathbb{B}$ such that $Member(s, \sigma)$ iff $s \in \lceil \sigma \rceil$.

The tuple $\mathcal{R}_S = (P, Pre, And, Diff, Empty)$ is called the *region algebra* of S . \square

Remark: Duality We take an existential view of symbolic transition systems. The dual, universal view requires (1) $\bigcap \{\lceil p \rceil \mid p \in P\} = \emptyset$, (2–4) closure of R under computable functions \overline{Pre} , \overline{And} , and \overline{Diff} such that

$$\lceil \overline{Pre}(\sigma) \rceil = \{s \in Q \mid (\forall t \in \delta(s) : t \in \sigma)\},$$

$\lceil \overline{And}(\sigma, \tau) \rceil = \lceil \sigma \rceil \cup \lceil \tau \rceil$, and $\lceil \overline{Diff}(\sigma, \tau) \rceil = Q \setminus \lceil Diff(\tau, \sigma) \rceil$, and (5) a computable function \overline{Empty} for deciding all universality questions about regions (that is, $\overline{Empty}(\sigma)$ iff $\lceil \sigma \rceil = Q$). All results of this paper have an alternative, dual formulation. \square

0.2 Example: Polyhedral Hybrid Automata

A *polyhedral hybrid automaton* H of dimension m , for a positive integer m , consists of the following components [AHH96]:

Continuous variables A set $X = \{x_1, \dots, x_m\}$ of real-valued variables. We write \dot{X} for the set $\{\dot{x}_1, \dots, \dot{x}_m\}$ of dotted variables (which represent first derivatives during continuous change), and we write X' for the set $\{x'_1, \dots, x'_m\}$ of primed variables (which represent values at the conclusion of discrete change). A *linear constraint* over X is an expression of the form $k_0 \sim k_1 x_1 + \dots + k_m x_m$, where $\sim \in \{<, \leq, =, \geq, >\}$ and k_0, \dots, k_m are integer constants. A *linear predicate* over X is a boolean combination of linear constraints over X . Let L^m be the set of linear predicates over X .

Discrete locations A finite directed multigraph (V, E) . The vertices in V are called *locations*; the edges in E are called *jumps*.

Invariant and flow conditions Two vertex-labeling functions *inv* and *flow*. For each location $v \in V$, the invariant condition *inv*(v) is a conjunction of linear constraints over X , and the flow condition *flow*(v) is a conjunction of linear constraints over \dot{X} . While the automaton control resides in location v , the variables may evolve according to *flow*(v) as long as *inv*(v) remains true.

Update conditions An edge-labeling functions *update*. For each jump $e \in E$, the update condition *update*(e) is a conjunction of linear constraints over $X \cup X'$. The predicate *update*(e) relates the possible values of the variables at the beginning of the jump (represented by X) and at the conclusion of the jump (represented by X').

The polyhedral hybrid automaton H is a *rectangular automaton* [HKPV98] if

- all linear constraints that occur in invariant conditions of H have the form $x \sim k$, for $x \in X$ and $k \in \mathbb{Z}$;

- all linear constraints that occur in flow conditions of H have the form $\dot{x} \sim k$, for $x \in X$ and $k \in \mathbb{Z}$;
- all linear constraints that occur in jump conditions of H have the form $x \sim k$ or $x' = x$ or $x' \sim k$, for $x \in X$ and $k \in \mathbb{Z}$;
- if e is a jump from location v to location v' , and $\text{update}(e)$ contains the conjunct $x' = x$, then both $\text{flow}(v)$ and $\text{flow}(v')$ contain the same constraints on \dot{x} .

The rectangular automaton H is a *singular automaton* if each flow condition of H has the form $\dot{x}_1 = k_1 \wedge \dots \wedge \dot{x}_m = k_m$. The singular automaton H is a *timed automaton* [AD94] if each flow condition of H has the form $\dot{x}_1 = 1 \wedge \dots \wedge \dot{x}_m = 1$.

The polyhedral hybrid automaton H defines the symbolic transition system $\mathcal{S}_H = (Q_H, \delta_H, R_H, \ulcorner \cdot \urcorner_H, P_H)$ with the following components:

- $Q_H = V \times \mathbb{R}^m$; that is, every state (v, \mathbf{x}) consists of a location v (the discrete component of the state) and values \mathbf{x} for the variables in X (the continuous component).
- $(v', \mathbf{x}') \in \delta_H(v, \mathbf{x})$ if either (1) there is a jump $e \in E$ from v to v' such that the closed predicate $\text{update}(e)[X, X' := \mathbf{x}, \mathbf{x}']$ is true, or (2) $v' = v$ and there is a real $\Delta \geq 0$ and a differentiable function $f: [0, \Delta] \rightarrow \mathbb{R}^m$ with first derivative \dot{f} such that $f(0) = \mathbf{x}$ and $f(\Delta) = \mathbf{x}'$, and for all reals $\varepsilon \in (0, \Delta)$, the closed predicates $\text{inv}(v)[X := f(\varepsilon)]$ and $\text{flow}(v)[\dot{X} := \dot{f}(\varepsilon)]$ are true. In case (2), the function f is called a *flow function*.
- $R_H = V \times L^m$; that is, every region (v, ϕ) consists of a location v (the discrete component of the region) and a linear predicate ϕ over X (the continuous component).
- $\ulcorner (v, \phi) \urcorner_H = \{(v, \mathbf{x}) \mid \mathbf{x} \in \mathbb{R}^m \text{ and } \phi[X := \mathbf{x}] \text{ is true}\}$; that is, the extension function maps the continuous component ϕ of a region to the values for the variables in X which satisfy the predicate ϕ . Consequently, the extension of every region consists of a location and a polyhedral subset of \mathbb{R}^m .
- $P_H = V \times \{\text{true}\}$; that is, only the discrete component of a state is observable.

It requires some work to see that \mathcal{S}_H is indeed a symbolic transition system. First, notice that the linear predicates over X are closed under all boolean operations, and that satisfiability is decidable for the linear predicates. Second, the *Pre* operator is computable on R_H , because all flow functions can be replaced by straight lines [AHH96].

0.3 Background Definitions

The symbolic transition systems are a special case of transition systems. A *transition system* $\mathcal{S} = (Q, \delta, \cdot, \ulcorner \cdot \urcorner, P)$ has the same components as a symbolic transition system, except that no regions are specified and the extension function is defined only for the observables (that is, $\ulcorner \cdot \urcorner: P \rightarrow 2^Q$).

State equivalences A *state equivalence* \cong is a family of relations which contains for each transition system \mathcal{S} an equivalence relation $\cong^{\mathcal{S}}$ on the states of \mathcal{S} .

The \cong *equivalence problem* for a class \mathbf{C} of transition systems asks, given two states s and t of a transition system \mathcal{S} from the class \mathbf{C} , whether $s \cong^{\mathcal{S}} t$. The state equivalence \cong_a is *as coarse as* the state equivalence \cong_b if $s \cong_a^{\mathcal{S}} t$ implies $s \cong_b^{\mathcal{S}} t$ for all transition systems \mathcal{S} . The equivalence \cong_a is *coarser than* \cong_b if \cong_a is as coarse as \cong_b , but \cong_b is not as coarse as \cong_a . Given a transition system $\mathcal{S} = (Q, \delta, \cdot, \ulcorner \cdot \urcorner, P)$ and a state equivalence \cong , the *quotient system* is the transition system $\mathcal{S}/\cong = (Q/\cong, \delta/\cong, \cdot, \ulcorner \cdot \urcorner/\cong, P)$ with the following components:

- the states in \mathcal{S}/\cong are the equivalence classes of $\cong_{\mathcal{S}}$;
- $\tau \in \delta/\cong(\sigma)$ if there is a state $s \in \sigma$ and a state $t \in \tau$ such that $t \in \delta(s)$;
- $\sigma \in \ulcorner p \urcorner/\cong$ if there is a state $s \in \sigma$ such that $s \in \ulcorner p \urcorner$.

The quotient construction is of particular interest to us when it transforms an infinite-state system \mathcal{S} into a finite-state system \mathcal{S}/\cong .

State logics A *state logic* L is a logic whose formulas are interpreted over the states of transition systems; that is, for every L -formula φ and every transition system \mathcal{S} , there is a set $\llbracket \varphi \rrbracket_{\mathcal{S}}$ of states of \mathcal{S} which satisfy φ . The *L model-checking problem* for a class \mathbf{C} of transition systems asks, given an L -formula φ and a state s of a transition system \mathcal{S} from the class \mathbf{C} , whether $s \in \llbracket \varphi \rrbracket_{\mathcal{S}}$. Two formulas φ and ψ of state logics are *equivalent* if $\llbracket \varphi \rrbracket_{\mathcal{S}} = \llbracket \psi \rrbracket_{\mathcal{S}}$ for all transition systems \mathcal{S} . The state logic L_a is *as expressive as* the state logic L_b if for every L_b -formula φ , there is an L_a -formula ψ which is equivalent to φ . The logic L_a is *more expressive than* L_b if L_a is as expressive as L_b , but L_b is not as expressive as L_a . Every state logic L *induces* a state equivalence, denoted \cong_L : for all states s and t of a transition system \mathcal{S} , define $s \cong_L^{\mathcal{S}} t$ if for all L -formulas φ , we have $s \in \llbracket \varphi \rrbracket_{\mathcal{S}}$ iff $t \in \llbracket \varphi \rrbracket_{\mathcal{S}}$. The state logic L *admits abstraction* if for every L -formula φ and every transition system \mathcal{S} , we have $\llbracket \varphi \rrbracket_{\mathcal{S}} = \bigcup \{ \sigma \mid \sigma \in \llbracket \varphi \rrbracket_{\mathcal{S}/\cong_L} \}$; that is, a state s of \mathcal{S} satisfies an L -formula φ iff the \cong_L equivalence class of s satisfies φ in the quotient system. Consequently, if L admits abstraction, then every L model-checking question on a transition system \mathcal{S} can be reduced to an L model-checking question on the induced quotient system \mathcal{S}/\cong_L . Below, we shall repeatedly prove the L model-checking problem for a class \mathbf{C} to be decidable by observing that for every transition system \mathcal{S} from \mathbf{C} , the quotient system \mathcal{S}/\cong_L has finitely many states and can be constructed effectively.

Symbolic semi-algorithms A *symbolic semi-algorithm* takes as input the region algebra $\mathcal{R}_{\mathcal{S}} = (P, Pre, And, Diff, Empty)$ of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \ulcorner \cdot \urcorner, P)$, and generates regions in R using the operations P , Pre , And , $Diff$, and $Empty$. Depending on the input \mathcal{S} , a symbolic semi-algorithm on \mathcal{S} may or may not terminate.

0.4 Preview

In sections 1–5 of this paper, we shall define five increasingly comprehensive classes of symbolic transition systems. In each case $i \in \{1, \dots, 5\}$, we will proceed in four steps:

1 Definition: Finite characterization We give a state equivalence \cong_i and define the class $\text{STS}(i)$ to contain precisely the symbolic transition systems \mathcal{S} for which the equivalence relation $\cong_i^{\mathcal{S}}$ has finite index (i.e., there are finitely many $\cong_i^{\mathcal{S}}$ equivalence classes). Each state equivalence \cong_i is coarser than its predecessor \cong_{i-1} , which implies that $\text{STS}(i-1) \subsetneq \text{STS}(i)$ for $i \in \{2, \dots, 5\}$.

2 Algorithmics: Symbolic state-space exploration We give a symbolic semi-algorithm that terminates precisely on the symbolic transition systems in the class $\text{STS}(i)$. This provides an operational characterization of the class $\text{STS}(i)$ which is equivalent to the denotational definition of $\text{STS}(i)$. Termination of the semi-algorithm is proved by observing that if given the region algebra of a symbolic transition system \mathcal{S} as input, then the extensions of all regions generated by the semi-algorithm are $\cong_i^{\mathcal{S}}$ blocks (i.e., unions of $\cong_i^{\mathcal{S}}$ equivalence classes). If \mathcal{S} is in the class $\text{STS}(i)$, then there are only finitely many $\cong_i^{\mathcal{S}}$ blocks, and the semi-algorithm terminates upon having constructed a representation of the quotient system \mathcal{S}/\cong_i . The semi-algorithm can therefore be used to decide all \cong_i equivalence questions for the class $\text{STS}(i)$.

3 Verification: Decidable properties We give a state logic L_i which admits abstraction and induces the state equivalence \cong_i . Since \cong_i quotients can be constructed effectively, it follows that the L_i model-checking problem for the class $\text{STS}(i)$ is decidable. However, model-checking algorithms which rely on the explicit construction of quotient systems are usually impractical. Hence, we also give a symbolic semi-algorithm that terminates on the symbolic transition systems in the class $\text{STS}(i)$ and directly decides all L_i model-checking questions for this class.

4 Example: Hybrid systems The interesting members of the class $\text{STS}(i)$ are those with infinitely many states. In four out of the five cases, following [Hen96], we provide certain kinds of polyhedral hybrid automata as examples.

1 Class-1 Symbolic Transition Systems

Class-1 systems are characterized by finite bisimilarity quotients. The region algebra of a class-1 system has a finite subalgebra that contains the observables and is closed under *Pre*, *And*, and *Diff* operations. This enables the model checking of all μ -calculus properties. Infinite-state examples of class-1 systems are provided by the singular hybrid automata.

1.1 Finite Characterization: Bisimilarity

Definition: Bisimilarity Let $\mathcal{S} = (Q, \delta, \cdot, \lceil \cdot \rceil, P)$ be a transition system. A binary relation \preceq on the state space Q is a *simulation* on \mathcal{S} if $s \preceq t$ implies the following two conditions:

1. For each observable $p \in P$, we have $s \in \lceil p \rceil$ iff $t \in \lceil p \rceil$.
2. For each state $s' \in \delta(s)$, there is a state $t' \in \delta(t)$ such that $s' \preceq t'$.

Symbolic semi-algorithm Closure1

Input: a region algebra $\mathcal{R} = (P, Pre, And, Diff, Empty)$.

```

 $T_0 := P;$ 
for  $i = 0, 1, 2, \dots$  do
   $T_{i+1} := T_i$ 
     $\cup \{Pre(\sigma) \mid \sigma \in T_i\}$ 
     $\cup \{And(\sigma, \tau) \mid \sigma, \tau \in T_i\}$ 
     $\cup \{Diff(\sigma, \tau) \mid \sigma, \tau \in T_i\}$ 
until  $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$ .

```

The termination test $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$, which is shorthand for $\{\lceil \sigma \rceil \mid \sigma \in T_{i+1}\} \subseteq \{\lceil \sigma \rceil \mid \sigma \in T_i\}$, is decided as follows: for each region $\sigma \in T_{i+1}$ check that there is a region $\tau \in T_i$ such that both $Empty(Diff(\sigma, \tau))$ and $Empty(Diff(\tau, \sigma))$.

Fig. 1. Partition refinement

Two states $s, t \in Q$ are *bisimilar*, denoted $s \cong_1^S t$, if there is a symmetric simulation \preceq on \mathcal{S} such that $s \preceq t$. The state equivalence \cong_1 is called *bisimilarity*. \square

Definition: Class STS1 A symbolic transition system \mathcal{S} belongs to the class STS1 if the bisimilarity relation \cong_1^S has finite index. \square

1.2 Symbolic State-space Exploration: Partition Refinement

The bisimilarity relation of a finite-state system can be computed by partition refinement [KS90]. The symbolic semi-algorithm Closure1 of Figure 1 applies this method to infinite-state systems [BFH90, Hen95]. Suppose that the input given to Closure1 is the region algebra of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \lceil \cdot \rceil, P)$. Then each T_i , for $i \geq 0$, is a finite set of regions; that is, $T_i \subseteq R$. By induction it is easy to check that for all $i \geq 0$, the extension of every region in T_i is a \cong_1^S block. Thus, if \cong_1^S has finite index, then Closure1 terminates. Conversely, suppose that Closure1 terminates with $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$. From the definition of bisimilarity it follows that if for each region $\sigma \in T_i$, we have $s \in \lceil \sigma \rceil$ iff $t \in \lceil \sigma \rceil$, then $s \cong_1^S t$. This implies that \cong_1^S has finite index.

Theorem 1A For all symbolic transition systems \mathcal{S} , the symbolic semi-algorithm Closure1 terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$ iff \mathcal{S} belongs to the class STS1.

Corollary 1A The \cong_1 (bisimilarity) equivalence problem is decidable for the class STS1 of symbolic transition systems.

1.3 Decidable Properties: Branching Time

Definition: μ -calculus The formulas of the μ -calculus are generated by the grammar

$$\varphi ::= p \mid \bar{p} \mid x \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \exists \bigcirc \varphi \mid \forall \bigcirc \varphi \mid (\mu x: \varphi) \mid (\nu x: \varphi),$$

for constants p from some set Π , and variables x from some set X . Let $\mathcal{S} = (Q, \delta, \cdot, \lceil \cdot \rceil, P)$ be a transition system whose observables include all constants; that is, $\Pi \subseteq P$. Let $\mathcal{E} : X \rightarrow 2^Q$ be a mapping from the variables to sets of states. We write $\mathcal{E}[x \mapsto \rho]$ for the mapping that agrees with \mathcal{E} on all variables, except that $x \in X$ is mapped to $\rho \subseteq Q$. Given \mathcal{S} and \mathcal{E} , every formula φ of the μ -calculus defines a set $\llbracket \varphi \rrbracket_{\mathcal{S}, \mathcal{E}} \subseteq Q$ of states:

$$\begin{aligned} \llbracket p \rrbracket_{\mathcal{S}, \mathcal{E}} &= \lceil p \rceil; \\ \llbracket \overline{p} \rrbracket_{\mathcal{S}, \mathcal{E}} &= Q \setminus \lceil p \rceil; \\ \llbracket x \rrbracket_{\mathcal{S}, \mathcal{E}} &= \mathcal{E}(x); \\ \llbracket \varphi_1 \{ \bigvee \} \varphi_2 \rrbracket_{\mathcal{S}, \mathcal{E}} &= \llbracket \varphi_1 \rrbracket_{\mathcal{S}, \mathcal{E}} \{ \bigcup \} \llbracket \varphi_2 \rrbracket_{\mathcal{S}, \mathcal{E}}; \\ \llbracket \{ \bigvee \} \bigcirc \varphi \rrbracket_{\mathcal{S}, \mathcal{E}} &= \{ s \in Q \mid (\{ \bigvee \} t \in \delta(s) : t \in \llbracket \varphi \rrbracket_{\mathcal{S}, \mathcal{E}}) \}; \\ \llbracket \{ \bigvee \} x : \varphi \rrbracket_{\mathcal{S}, \mathcal{E}} &= \{ \bigcup \} \{ \rho \subseteq Q \mid \rho = \llbracket \varphi \rrbracket_{\mathcal{S}, \mathcal{E}[x \mapsto \rho]} \}. \end{aligned}$$

If we restrict ourselves to the closed formulas of the μ -calculus, then we obtain a state logic, denoted L_1^μ : the state $s \in Q$ satisfies the L_1^μ -formula φ if $s \in \llbracket \varphi \rrbracket_{\mathcal{S}, \mathcal{E}}$ for any variable mapping \mathcal{E} ; that is, $\llbracket \varphi \rrbracket_{\mathcal{S}} = \llbracket \varphi \rrbracket_{\mathcal{S}, \mathcal{E}}$ for any \mathcal{E} . \square

Remark: Duality For every L_1^μ -formula φ , the dual L_1^μ -formula $\overline{\varphi}$ is obtained by replacing the constructors $p, \overline{p}, \vee, \wedge, \exists \bigcirc, \forall \bigcirc, \mu$, and ν by $\overline{p}, p, \wedge, \vee, \forall \bigcirc, \exists \bigcirc, \nu$, and μ , respectively. Then, $\llbracket \overline{\varphi} \rrbracket_{\mathcal{S}} = Q \setminus \llbracket \varphi \rrbracket_{\mathcal{S}}$. It follows that the answer of the model-checking question for a state $s \in Q$ and an L_1^μ -formula φ is complementary to the answer of the model-checking question for s and the dual formula $\overline{\varphi}$. \square

The following facts about the μ -calculus are relevant in our context [AH98]. First, L_1^μ admits abstraction, and the state equivalence induced by L_1^μ is \cong_1 (bisimilarity). Second, L_1^μ is very expressive; in particular, L_1^μ is more expressive than the temporal logics CTL^* and CTL , which also induce bisimilarity. Third, the definition of L_1^μ naturally suggests a model-checking method for finite-state systems, where each fixpoint can be computed by successive approximation. The symbolic semi-algorithm **ModelCheck** of Figure 2 applies this method to infinite-state systems.

Suppose that the input given to **ModelCheck** is the region algebra of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \lceil \cdot \rceil, P)$, a μ -calculus formula φ , and any mapping $E : X \rightarrow 2^R$ from the variables to sets of regions. Then for each recursive call of **ModelCheck**, each T_i , for $i \geq 0$, is a finite set of regions from R , and each recursive call returns a finite set of regions from R . It is easy to check that all of these regions are also generated by the semi-algorithm **Closure1** on input $\mathcal{R}_{\mathcal{S}}$. Thus, if **Closure1** terminates, then so does **ModelCheck**. Furthermore, if it terminates, then **ModelCheck** returns a set $[\varphi]_E \subseteq R$ of regions such that $\bigcup \{ \lceil \sigma \rceil \mid \sigma \in [\varphi]_E \} = \llbracket \varphi \rrbracket_{\mathcal{S}, \mathcal{E}}$, where $\mathcal{E}(x) = \bigcup \{ \lceil \sigma \rceil \mid \sigma \in E(x) \}$ for all $x \in X$. In particular, if φ is closed, then a state $s \in Q$ satisfies φ iff $\text{Member}(s, \sigma)$ for some region $\sigma \in [\varphi]_E$.

Theorem 1B. *For all symbolic transition systems \mathcal{S} in STS1 and every L_1^μ -formula φ , the symbolic semi-algorithm **ModelCheck** terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$ and the input formula φ .*

Symbolic semi-algorithm ModelCheck

Input: a region algebra $\mathcal{R} = (P, Pre, And, Diff, Empty)$, a formula $\varphi \in L_1^\mu$, and a mapping E with domain X .

Output: $[\varphi]_E :=$

```

if  $\varphi = p$  then return  $\{p\}$ ;
if  $\varphi = \bar{p}$  then return  $\{Diff(q, p) \mid q \in P\}$ ;
if  $\varphi = (\varphi_1 \vee \varphi_2)$  then return  $[\varphi_1]_E \cup [\varphi_2]_E$ ;
if  $\varphi = (\varphi_1 \wedge \varphi_2)$  then
    return  $\{And(\sigma, \tau) \mid \sigma \in [\varphi_1]_E \text{ and } \tau \in [\varphi_2]_E\}$ ;
if  $\varphi = \exists \bigcirc \varphi'$  then return  $\{Pre(\sigma) \mid \sigma \in [\varphi']_E\}$ ;
if  $\varphi = \forall \bigcirc \varphi'$  then return  $P \setminus \setminus \{Pre(\sigma) \mid \sigma \in (P \setminus \setminus [\varphi']_E)\}$ ;
if  $\varphi = (\mu x: \varphi')$  then
     $T_0 := \emptyset$ ;
    for  $i = 0, 1, 2, \dots$  do
         $T_{i+1} := [\varphi']_{E[x \mapsto T_i]}$ 
        until  $\bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_{i+1}\} \subseteq \bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_i\}$ ;
    return  $T_i$ ;
if  $\varphi = (\nu x: \varphi')$  then
     $T_0 := P$ ;
    for  $i = 0, 1, 2, \dots$  do
         $T_{i+1} := [\varphi']_{E[x \mapsto T_i]}$ 
        until  $\bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_{i+1}\} \supseteq \bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_i\}$ ;
    return  $T_i$ .

```

The *pairwise-difference* operation $T \setminus \setminus T'$ between two finite sets T and T' of regions is computed inductively as follows:

$$\begin{aligned}
 T \setminus \setminus \emptyset &= T; \\
 T \setminus \setminus (\{\tau\} \cup T') &= \{Diff(\sigma, \tau) \mid \sigma \in T\} \setminus \setminus T'.
 \end{aligned}$$

The termination test $\bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T\} \subseteq \bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T'\}$ is decided by checking that $Empty(\sigma)$ for each region $\sigma \in (T \setminus \setminus T')$.

Fig. 2. Model checking

Corollary 1B *The L_1^μ model-checking problem is decidable for the class STS1 of symbolic transition systems.*

1.4 Example: Singular Hybrid Automata

The fundamental theorem of timed automata [AD94] shows that for every timed automaton, the (time-abstract) bisimilarity relation has finite index. The proof can be extended to the singular automata [ACH⁺95]. It follows that the symbolic semi-algorithm **ModelCheck**, which has been implemented for polyhedral hybrid automata in the tool **HyTECH** [HHWT95], decides all L_1^μ model-checking questions for singular automata. The singular automata form a maximal class

of hybrid automata in STS1. This is because there is a 2D (two-dimensional) rectangular automaton whose bisimilarity relation is state equality [Hen95].

Theorem 1C *The singular automata belong to the class STS1. There is a 2D rectangular automaton that does not belong to STS1.*

2 Class-2 Symbolic Transition Systems

Class-2 systems are characterized by finite similarity quotients. The region algebra of a class-2 system has a finite subalgebra that contains the observables and is closed under *Pre* and *And* operations. This enables the model checking of all existential and universal μ -calculus properties. Infinite-state examples of class-2 systems are provided by the 2D rectangular hybrid automata.

2.1 Finite Characterization: Similarity

Definition: Similarity Let \mathcal{S} be a transition system. Two states s and t of \mathcal{S} are *similar*, denoted $s \cong_2^S t$, if there is a simulation \preceq on \mathcal{S} such that both $s \preceq t$ and $t \preceq s$. The state equivalence \cong_2 is called *similarity*. \square

Definition: Class STS2 A symbolic transition system \mathcal{S} belongs to the class STS2 if the similarity relation \cong_2^S has finite index. \square

Since similarity is coarser than bisimilarity [vG90], the class STS2 of symbolic transition systems is a proper extension of STS1.

2.2 Symbolic State-space Exploration: Intersection Refinement

The symbolic semi-algorithm **Closure2** of Figure 3 is an abstract version of the method presented in [HHK95] for computing the similarity relation of an infinite-state system. Suppose that the input given to **Closure2** is the region algebra of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \ulcorner \cdot \urcorner, P)$. Given two states $s, t \in Q$, we say that t *simulates* s if $s \preceq t$ for some simulation \preceq on \mathcal{S} . For $i \geq 0$ and $s \in Q$, define

$$Sim_i(s) = \bigcap \{ \ulcorner \sigma \urcorner \mid \sigma \in T_i \text{ and } s \in \ulcorner \sigma \urcorner \},$$

where the set T_i of regions is computed by **Closure2**. By induction it is easy to check that for all $i \geq 0$, if t simulates s , then $t \in Sim_i(s)$. Thus, the extension of every region in T_i is a \cong_2^S block, and if \cong_2^S has finite index, then **Closure2** terminates. Conversely, suppose that **Closure2** terminates with $\ulcorner T_{i+1} \urcorner \subseteq \ulcorner T_i \urcorner$. From the definition of simulations it follows that if $t \in Sim_i(s)$, then t simulates s . This implies that \cong_2^S has finite index.

Theorem 2A *For all symbolic transition systems \mathcal{S} , the symbolic semi-algorithm **Closure2** terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$ iff \mathcal{S} belongs to the class STS2.*

Corollary 2A *The \cong_2 (similarity) equivalence problem is decidable for the class STS2 of symbolic transition systems.*

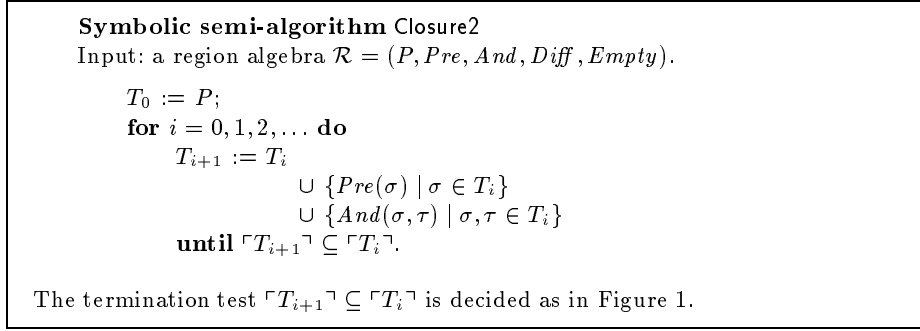


Fig. 3. Intersection refinement

2.3 Decidable Properties: Negation-free Branching Time

Definition: Negation-free μ -calculus The *negation-free μ -calculus* consists of the μ -calculus formulas that are generated by the grammar

$$\varphi ::= p \mid x \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \exists \bigcirc \varphi \mid (\mu x: \varphi) \mid (\nu x: \varphi),$$

for constants $p \in \Pi$ and variables $x \in X$. The state logic L_2^μ consists of the closed formulas of the negation-free μ -calculus. The state logic \overline{L}_2^μ consists of the duals of all L_2^μ -formulas. \square

The following facts about the negation-free μ -calculus and its dual are relevant in our context [AH98]. First, both L_2^μ and \overline{L}_2^μ admit abstraction, and the state equivalence induced by both L_2^μ and \overline{L}_2^μ is \cong_2 (similarity). It follows that the logic L_1^μ with negation is more expressive than either L_2^μ or \overline{L}_2^μ . Second, the negation-free logic L_2^μ is more expressive than the existential fragments of CTL* and CTL, which also induce similarity, and the dual logic \overline{L}_2^μ is more expressive than the universal fragments of CTL* and CTL, which again induce similarity.

If we apply the symbolic semi-algorithm **ModelCheck** of Figure 2 to the region algebra of a symbolic transition system \mathcal{S} and an input formula from L_2^μ , then the cases $\varphi = \overline{p}$ and $\varphi = \forall \bigcirc \varphi'$ are never executed. It follows that all regions which are generated by **ModelCheck** are also generated by the semi-algorithm **Closure2** on input $\mathcal{R}_\mathcal{S}$. Thus, if **Closure2** terminates, then so does **ModelCheck**.

Theorem 2B *For all symbolic transition systems \mathcal{S} in STS2 and every L_2^μ -formula φ , the symbolic semi-algorithm **ModelCheck** terminates on the region algebra $\mathcal{R}_\mathcal{S}$ and the input formula φ .*

Corollary 2B *The L_2^μ and \overline{L}_2^μ model-checking problems are decidable for the class STS2 of symbolic transition systems.*

2.4 Example: 2D Rectangular Hybrid Automata

For every 2D rectangular automaton, the (time-abstract) similarity relation has finite index [HHK95]. It follows that the symbolic semi-algorithm **ModelCheck**, as implemented in **HYTECH**, decides all L_2^μ and $\overline{L_2^\mu}$ model-checking questions for 2D rectangular automata. The 2D rectangular automata form a maximal class of hybrid automata in STS2. This is because there is a 3D rectangular automaton whose similarity relation is state equality [HK96].

Theorem 2C *The 2D rectangular automata belong to the class STS2. There is a 3D rectangular automaton that does not belong to STS2.*

3 Class-3 Symbolic Transition Systems

Class-3 systems are characterized by finite trace-equivalence quotients. The region algebra of a class-3 system has a finite subalgebra that contains the observables and is closed under *Pre* operations and those *And* operations for which one of the two arguments is an observable. This enables the model checking of all linear temporal properties. Infinite-state examples of class-3 systems are provided by the rectangular hybrid automata.

3.1 Finite Characterization: Traces

Definition: Trace equivalence Let $\mathcal{S} = (Q, \delta, \cdot, \lceil \cdot \rceil, P)$ be a transition system. Given a state $s_0 \in Q$, a *source- s_0 trace* π of \mathcal{S} is a finite sequence $p_0 p_1 \dots p_n$ of observables $p_i \in P$ such that

1. $s_0 \in \lceil p_0 \rceil$;
2. for all $0 \leq i < n$, there is a state $s_{i+1} \in (\delta(s_i) \cap \lceil p_{i+1} \rceil)$.

The number n of observables (minus 1) is called the *length* of the trace π , the final state s_n is the *sink* of π , and the final observable p_n is the *target* of π . Two states $s, t \in Q$ are *trace equivalent*, denoted $s \cong_3^{\mathcal{S}} t$, if every source- s trace of \mathcal{S} is a source- t trace of \mathcal{S} , and vice versa. The state equivalence \cong_3 is called *trace equivalence*. \square

Definition: Class STS3 A symbolic transition system \mathcal{S} belongs to the class STS3 if the trace-equivalence relation $\cong_3^{\mathcal{S}}$ has finite index. \square

Since trace equivalence is coarser than similarity [vG90], the class STS3 of symbolic transition systems is a proper extension of STS2.

3.2 Symbolic State-space Exploration: Observation Refinement

Trace equivalence can be characterized operationally by the symbolic semi-algorithm **Closure3** of Figure 4. We shall show that, when the input is the region

Symbolic semi-algorithm Closure3
Input: a region algebra $\mathcal{R} = (P, Pre, And, Diff, Empty)$.

```

 $T_0 := P;$ 
for  $i = 0, 1, 2, \dots$  do
   $T_{i+1} := T_i$ 
     $\cup \{Pre(\sigma) \mid \sigma \in T_i\}$ 
     $\cup \{And(\sigma, p) \mid \sigma \in T_i \text{ and } p \in P\}$ 
  until  $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$ .

```

The termination test $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$ is decided as in Figure 1.

Fig. 4. Observation refinement

algebra of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \lceil \cdot \rceil, P)$, then **Closure3** terminates iff the trace-equivalence relation $\cong_3^{\mathcal{S}}$ has finite index. Furthermore, upon termination, $s \cong_3^{\mathcal{S}} t$ iff for each region $\sigma \in T_i$, we have $s \in \lceil \sigma \rceil$ iff $t \in \lceil \sigma \rceil$.

Theorem 3A *For all symbolic transition systems \mathcal{S} , the symbolic semi-algorithm Closure3 terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$ iff \mathcal{S} belongs to the class STS3.*

Proof [HM99] We proceed in two steps. First, we show that **Closure3** terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$ iff the equivalence relation $\cong_{L_3^\mu}^{\mathcal{S}}$ induced by the linear-time μ -calculus (defined below) has finite index. Second, we show that $\cong_{L_3^\mu}^{\mathcal{S}}$ coincides with trace equivalence. The proof of the first part proceeds as usual. It can be seen by induction that for all $i \geq 0$, the extension of every region in T_i , as computed by **Closure3**, is a $\cong_{L_3^\mu}^{\mathcal{S}}$ block. Thus, if $\cong_{L_3^\mu}^{\mathcal{S}}$ has finite index, then **Closure3** terminates. Conversely, suppose that **Closure3** terminates with $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$. It can be shown that if two states are not $\cong_{L_3^\mu}^{\mathcal{S}}$ -equivalent, then there is a region in T_i which contains one state but not the other. It follows that if for each region $\sigma \in T_i$, we have $s \in \lceil \sigma \rceil$ iff $t \in \lceil \sigma \rceil$, then $s \cong_{L_3^\mu}^{\mathcal{S}} t$. This implies that $\cong_{L_3^\mu}^{\mathcal{S}}$ has finite index.

For the second part, we show that L_3^μ is as expressive as the logic $\exists B\ddot{U}CHI$, whose formulas are the existentially interpreted Büchi automata, and that $\exists B\ddot{U}CHI$ is as expressive as L_3^μ . This result is implicit in a proof by [EJS93]. By induction on the structure of an L_3^μ -formula φ , we can construct a Büchi automaton B_φ such that for all transition systems \mathcal{S} , a state s of \mathcal{S} satisfies φ iff for some infinite source- s trace of \mathcal{S} is accepted by B_φ . Conversely, given a Büchi automaton B , we can construct an L_3^μ -formula which is equivalent to $\exists B$ [Dam94]. Since the state equivalence induced by $\exists B\ddot{U}CHI$ is trace equivalence, it follows that $\cong_{L_3^\mu}^{\mathcal{S}}$ is also trace equivalence. \square

Corollary 3A *The \cong_3 (trace) equivalence problem is decidable for the class STS3 of symbolic transition systems.*

3.3 Decidable Properties: Linear Time

Definition: Linear-time μ -calculus The *linear-time μ -calculus* (also called “ L_1 ” in [EJS93]) consists of the μ -calculus formulas that are generated by the grammar

$$\varphi ::= p \mid x \mid \varphi \vee \varphi \mid p \wedge \varphi \mid \exists \bigcirc \varphi \mid (\mu x : \varphi) \mid (\nu x : \varphi),$$

for constants $p \in \Pi$ and variables $x \in X$. The state logic L_3^μ consists of the closed formulas of the linear-time μ -calculus. The state logic \overline{L}_3^μ consists of the duals of all L_3^μ -formulas. \square

The following facts about the linear-time μ -calculus and its dual are relevant in our context (cf. the second part of the proof of Theorem 3A). First, both L_3^μ and \overline{L}_3^μ admit abstraction, and the state equivalence induced by both L_3^μ and \overline{L}_3^μ is \cong_3 (trace equivalence). It follows that the logic L_2^μ with unrestricted conjunction is more expressive than L_3^μ , and \overline{L}_2^μ is more expressive than \overline{L}_3^μ . Second, the logic L_3^μ with restricted conjunction is more expressive than the existential interpretation of the linear temporal logic LTL, which also induces trace equivalence. For example, the existential LTL formula $\exists(p\mathcal{U}q)$ (“on some trace, p until q ”) is equivalent to the L_3^μ -formula $(\mu x : q \vee (p \wedge \exists \bigcirc x))$ (notice that one argument of the conjunction is a constant). The dual logic \overline{L}_3^μ is more expressive than the usual, universal interpretation of LTL, which again induces trace equivalence. For example, the (universal) LTL formula $p\mathcal{W}q$ (“on all traces, either p forever, or p until q ”) is equivalent to the \overline{L}_3^μ -formula $(\nu x : p \wedge \forall \bigcirc (q \vee x))$ (notice that one argument of the disjunction is a constant).

If we apply the symbolic semi-algorithm **ModelCheck** of Figure 2 to the region algebra of a symbolic transition system \mathcal{S} and an input formula from L_3^μ , then all regions which are generated by **ModelCheck** are also generated by the semi-algorithm **Closure3** on input $\mathcal{R}_\mathcal{S}$. Thus, if **Closure3** terminates, then so does **ModelCheck**.

Theorem 3B *For all symbolic transition systems \mathcal{S} in STS3 and every L_3^μ -formula φ , the symbolic semi-algorithm **ModelCheck** terminates on the region algebra $\mathcal{R}_\mathcal{S}$ and the input formula φ .*

Corollary 3B *The L_3^μ and \overline{L}_3^μ model-checking problems are decidable for the class STS3 of symbolic transition systems.*

Remark: LTL model checking These results suggest, in particular, a symbolic procedure for model checking LTL properties over STS3 systems [HM99]. Suppose that \mathcal{S} is a symbolic transition system in the class STS3, and φ is an LTL formula. First, convert $\neg\varphi$ to a Büchi automaton $B_{\neg\varphi}$ using a tableau construction, and then to an equivalent L_3^μ -formula ψ (introduce one variable per state of $B_{\neg\varphi}$). Second, run the symbolic semi-algorithm **ModelCheck** on inputs $\mathcal{R}_\mathcal{S}$ and ψ . It will terminate with a representation of the complement of the set of states that satisfy φ in \mathcal{S} . \square




Fig. 5. Distance equivalence is coarser than trace equivalence

3.4 Example: Rectangular Hybrid Automata

For every rectangular automaton, the (time-abstract) trace-equivalence relation has finite index [HKPV98]. It follows that the symbolic semi-algorithm **ModelCheck**, as implemented in **HYTECH**, decides all L_3^μ and \overline{L}_3^μ model-checking questions for rectangular automata. The rectangular automata form a maximal class of hybrid automata in **STS3**. This is because for simple generalizations of rectangular automata, the reachability problem is undecidable [HKPV98].

Theorem 3C *The rectangular automata belong to the class STS3.*

4 Class-4 Symbolic Transition Systems

We define two states of a transition system to be “distance equivalent” if for every distance d , the same observables can be reached in d transitions. Class-4 systems are characterized by finite distance-equivalence quotients. The region algebra of a class-4 system has a finite subalgebra that contains the observables and is closed under *Pre* operations. This enables the model checking of all existential conjunction-free and universal disjunction-free μ -calculus properties, such as the property that an observable can be reached in an even number of transitions.

4.1 Finite Characterization: Equi-distant Targets

Definition: Distance equivalence Let \mathcal{S} be a transition system. Two states s and t of \mathcal{S} are *distance equivalent*, denoted $s \cong_4^{\mathcal{S}} t$, if for every source- s trace of \mathcal{S} with length n and target p , there is a source- t trace of \mathcal{S} with length n and target p , and vice versa. The state equivalence \cong_4 is called *distance equivalence*. \square

Definition: Class STS4 A symbolic transition system \mathcal{S} belongs to the class **STS4** if the distance-equivalence relation $\cong_4^{\mathcal{S}}$ has finite index. \square

Figure 5 shows that distance equivalence is coarser than trace equivalence (s and t are distance equivalent but not trace equivalent). It follows that the class **STS4** of symbolic transition systems is a proper extension of **STS3**.

Symbolic semi-algorithm Closure4Input: a region algebra $\mathcal{R} = (P, Pre, \cdot, Diff, Empty)$.

```

 $T_0 := P;$ 
for  $i = 0, 1, 2, \dots$  do
   $T_{i+1} := T_i$ 
   $\cup \{Pre(\sigma) \mid \sigma \in T_i\}$ 
until  $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$ .

```

The termination test $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$ is decided as in Figure 1.**Fig. 6.** Predecessor iteration**4.2 Symbolic State-space Exploration: Predecessor Iteration**

The symbolic semi-algorithm **Closure4** of Figure 6 computes the subalgebra of a region algebra \mathcal{R}_S that contains the observables and is closed under the *Pre* operation. Suppose that the input given to **Closure4** is the region algebra of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \lceil \cdot \rceil, P)$. For $i \geq 0$ and $s, t \in Q$, define $s \sim_i^S t$ if for every source- s trace of \mathcal{S} with length $n \leq i$ and target p , there is a source- t trace of \mathcal{S} with length n and target p , and vice versa. By induction it is easy to check that for all $i \geq 0$, the extension of every region in T_i , as computed by **Closure4**, is a \sim_i^S block. Since \sim_i^S is as coarse as \sim_{i+1}^S for all $i \geq 0$, and \cong_2^S is equal to $\bigcap \{\sim_i^S \mid i \geq 0\}$, if \cong_2^S has finite index, then \cong_2^S is equal to \sim_i^S for some $i \geq 0$. Then, **Closure2** will terminate in i iterations. Conversely, suppose that **Closure4** terminates with $\lceil T_{i+1} \rceil \subseteq \lceil T_i \rceil$. In this case, if for all regions $\sigma \in T_i$, we have $s \in \lceil \sigma \rceil$ iff $t \in \lceil \sigma \rceil$, then $s \cong_4^S t$. This is because if s can reach an observable p in n transitions, but t cannot, then there is a region in T_i , namely, $Pre^n(p)$, such that $s \in \lceil Pre^n(p) \rceil$ and $t \notin \lceil Pre^n(p) \rceil$. It follows that \cong_4^S has finite index.

Theorem 4A *For all symbolic transition systems \mathcal{S} , the symbolic semi-algorithm Closure4 terminates on the region algebra \mathcal{R}_S iff \mathcal{S} belongs to the class STS4.*

Corollary 4A *The \cong_4 (distance) equivalence problem is decidable for the class STS4 of symbolic transition systems.*

4.3 Decidable Properties: Conjunction-free Linear Time

Definition: Conjunction-free μ -calculus The *conjunction-free μ -calculus* consists of the μ -calculus formulas that are generated by the grammar

$$\varphi ::= p \mid x \mid \varphi \vee \varphi \mid \exists \bigcirc \varphi \mid (\mu x : \varphi)$$

for constants $p \in \Pi$ and variables $x \in X$. The state logic L_4^μ consists of the closed formulas of the conjunction-free μ -calculus. The state logic \overline{L}_4^μ consists of the duals of all L_4^μ -formulas. \square

Definition: Conjunction-free temporal logic The formulas of the *conjunction-free temporal logic* L_4^\diamond are generated by the grammar

$$\varphi ::= p \mid \varphi \vee \varphi \mid \exists \bigcirc \varphi \mid \exists \diamond_{\leq d} \varphi \mid \exists \diamond \varphi,$$

for constants $p \in \Pi$ and nonnegative integers d . Let $\mathcal{S} = (Q, \delta, \cdot, \lceil \cdot \rceil, P)$ be a transition system whose observables include all constants; that is, $\Pi \subseteq P$. The L_4^\diamond -formula φ defines the set $\llbracket \varphi \rrbracket_{\mathcal{S}} \subseteq Q$ of satisfying states:

$$\begin{aligned} \llbracket p \rrbracket_{\mathcal{S}} &= \lceil p \rceil; \\ \llbracket \varphi_1 \vee \varphi_2 \rrbracket_{\mathcal{S}} &= \llbracket \varphi_1 \rrbracket_{\mathcal{S}} \cup \llbracket \varphi_2 \rrbracket_{\mathcal{S}}; \\ \llbracket \exists \bigcirc \varphi \rrbracket_{\mathcal{S}} &= \{s \in Q \mid (\exists t \in \delta(s): t \in \llbracket \varphi \rrbracket_{\mathcal{S}})\}; \\ \llbracket \exists \diamond_{\leq d} \varphi \rrbracket_{\mathcal{S}} &= \{s \in Q \mid \text{there is a source-}s \text{ trace of } \mathcal{S} \text{ with} \\ &\quad \text{length at most } d \text{ and sink in } \llbracket \varphi \rrbracket_{\mathcal{S}}\}; \\ \llbracket \exists \diamond \varphi \rrbracket_{\mathcal{S}} &= \{s \in Q \mid \text{there is a source-}s \text{ trace of } \mathcal{S} \text{ with sink in } \llbracket \varphi \rrbracket_{\mathcal{S}}\}. \end{aligned}$$

(The constructor $\exists \diamond_{\leq d}$ is definable from $\exists \bigcirc$ and \vee ; however, it will be essential in the $\exists \bigcirc$ -free fragment of L_4^\diamond we will consider below.) \square

Remark: Duality For every L_4^\diamond -formula φ , the *dual* formula $\overline{\varphi}$ is obtained by replacing the constructors p , \vee , $\exists \bigcirc$, $\exists \diamond_{\leq d}$, and $\exists \diamond$ by \overline{p} , \wedge , $\forall \bigcirc$, $\forall \square_{\leq d}$, and $\forall \square$, respectively. The semantics of the dual constructors is defined as usual, such that $\llbracket \overline{\varphi} \rrbracket_{\mathcal{S}} = Q \setminus \llbracket \varphi \rrbracket_{\mathcal{S}}$. The state logic $\overline{L_4^\diamond}$ consists of the duals of all L_4^\diamond -formulas. It follows that the answer of the model-checking question for a state $s \in Q$ and an $\overline{L_4^\diamond}$ -formula $\overline{\varphi}$ is complementary to the answer of the model-checking question for s and the L_4^\diamond -formula φ . \square

The following facts about the conjunction-free μ -calculus, conjunction-free temporal logic, and their duals are relevant in our context. First, both L_4^μ and $\overline{L_4^\mu}$ admit abstraction, and the state equivalence induced by both L_4^μ and $\overline{L_4^\mu}$ is \cong_4 (distance equivalence). It follows that the logic L_3^μ with restricted conjunction is more expressive than L_4^μ , and $\overline{L_3^\mu}$ is more expressive than $\overline{L_4^\mu}$. Second, the conjunction-free μ -calculus L_4^μ is more expressive than the conjunction-free temporal logic L_4^\diamond , and $\overline{L_4^\mu}$ is more expressive than $\overline{L_4^\diamond}$, both of which also induce distance equivalence. For example, the property that an observable can be reached in an even number of transitions can be expressed in L_4^μ but not in L_4^\diamond .

If we apply the symbolic semi-algorithm **ModelCheck** of Figure 2 to the region algebra of a symbolic transition system \mathcal{S} and an input formula from L_4^μ , then all regions which are generated by **ModelCheck** are also generated by the semi-algorithm **Closure4** on input $\mathcal{R}_{\mathcal{S}}$. Thus, if **Closure4** terminates, then so does **ModelCheck**.

Theorem 4B *For all symbolic transition systems \mathcal{S} in STS4 and every L_4^μ -formula φ , the symbolic semi-algorithm **ModelCheck** terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$ and the input formula φ .*

Corollary 4B *The L_4^μ and $\overline{L_4^\mu}$ model-checking problems are decidable for the class STS4 of symbolic transition systems.*




Fig. 7. Bounded-reach equivalence is coarser than distance equivalence

5 Class-5 Symbolic Transition Systems

We define two states of a transition system to be “bounded-reach equivalent” if for every distance d , the same observables can be reached in d or fewer transitions. Class-5 systems are characterized by finite bounded-reach-equivalence quotients. Equivalently, for every observable p there is a finite bound n_p such that all states that can reach p can do so in at most n_p transitions. This enables the model checking of all reachability and (by duality) invariance properties. The transition systems in class 5 have also been called “well-structured” [AČJT96]. Infinite-state examples of class-5 systems are provided by networks of rectangular hybrid automata.

5.1 Finite Characterization: Bounded-distance Targets

Definition: Bounded-reach equivalence Let \mathcal{S} be a transition system. Two states s and t of \mathcal{S} are *bounded-reach equivalent*, denoted $s \cong_5^{\mathcal{S}} t$, if for every source- s trace of \mathcal{S} with length n and target p , there is a source- t trace of \mathcal{S} with length at most n and target p , and vice versa. The state equivalence \cong_5 is called *bounded-reach equivalence*. \square

Definition: Class STS5 A symbolic transition system \mathcal{S} belongs to the class STS5 if the bounded-reach-equivalence relation $\cong_5^{\mathcal{S}}$ has finite index. \square

Figure 7 shows that bounded-reach equivalence is coarser than distance equivalence (all states s_i , for $i \geq 0$, are bounded-reach equivalent, but no two of them are distance equivalent). It follows that the class STS5 of symbolic transition systems is a proper extension of STS4.

5.2 Symbolic State-space Exploration: Predecessor Aggregation

The symbolic semi-algorithm **Reach** of Figure 8 starts from the observables and repeatedly applies the *Pre* operation, but its termination criterion is more easily met than the termination criterion of the semi-algorithm **Closure4**; that is, **Reach** may terminate on more inputs than **Closure4**. Indeed, we shall show

Symbolic semi-algorithm Reach
Input: a region algebra $\mathcal{R} = (P, Pre, And, Diff, Empty)$.

```

for each  $p \in P$  do
   $T_0 := \{p\};$ 
  for  $i = 0, 1, 2, \dots$  do
     $T_{i+1} := T_i \cup \{Pre(\sigma) \mid \sigma \in T_i\}$ 
    until  $\bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_{i+1}\} \subseteq \bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_i\}$ 
  end.

```

The termination test $\bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_{i+1}\} \subseteq \bigcup \{\ulcorner \sigma \urcorner \mid \sigma \in T_i\}$ is decided as in Figure 2.

Fig. 8. Predecessor aggregation

that, when the input is the region algebra of a symbolic transition system $\mathcal{S} = (Q, \delta, R, \ulcorner \cdot \urcorner, P)$, then **Reach** terminates iff \mathcal{S} belongs to the class **STS5**. Furthermore, upon termination, $s \cong_5^{\mathcal{S}} t$ iff for each observation $p \in P$ and each region $\sigma \in T_i^p$, we have $s \in \ulcorner \sigma \urcorner$ iff $t \in \ulcorner \sigma \urcorner$.

An alternative characterization of the class **STS5** can be given using well-quasi-orders on states [AČJT96, FS98]. A *quasi-order* on a set A is a reflexive and transitive binary relation on A . A *well-quasi-order* on A is a quasi-order \preceq on A such that for every infinite sequence a_0, a_1, a_2, \dots of elements $a_i \in A$ there exist indices i and j with $i < j$ and $a_i \preceq a_j$. A set $B \subseteq A$ is *upward-closed* if for all $b \in B$ and $a \in A$, if $b \preceq a$, then $a \in B$. It can be shown that if \preceq is a well-quasi-order on A , then every infinite increasing sequence $B_0 \subseteq B_1 \subseteq B_2 \subseteq \dots$ of upward-closed sets $B_i \subseteq A$ eventually stabilizes; that is, there exists an index $i \geq 0$ such that $B_j = B_i$ for all $j \geq i$.

Theorem 5A. *For all symbolic transition systems \mathcal{S} , the following three conditions are equivalent:*

1. \mathcal{S} belongs to the class **STS5**.
2. The symbolic semi-algorithm **Reach** terminates on the region algebra $\mathcal{R}_{\mathcal{S}}$.
3. There is a well-quasi-order \preceq on the states of \mathcal{S} such that for all observations p and all nonnegative integers d , the set $\llbracket \exists \Diamond_{\leq d} p \rrbracket_{\mathcal{S}}$ is upward-closed.

Proof ($2 \Rightarrow 1$) Define $s \sim_{\leq n}^{\mathcal{S}} t$ if for all observations p , for every source- s trace with length n and target p , there is a source- t trace with length at most n and target p , and vice versa. Note that $\sim_{\leq n}^{\mathcal{S}}$ has finite index for all $n \geq 0$. Suppose that the semi-algorithm **Reach** terminates in at most i iterations for each observation p . Then for all $n \geq i$, the equivalence relation $\sim_{\leq n}^{\mathcal{S}}$ is equal to $\sim_{\leq i}^{\mathcal{S}}$. Since $\cong_5^{\mathcal{S}}$ is equal to $\bigcap \{\sim_{\leq n}^{\mathcal{S}} \mid n \geq 0\}$, it has finite index.

($1 \Rightarrow 3$) Define the quasi-order $s \preceq_5^{\mathcal{S}} t$ if for all observables p and all $n \geq 0$, for every source- s trace with length n and target p , there is a source- t trace with

length at most n and target p . Then each set $\llbracket \exists \Diamond_{\leq d} p \rrbracket s$, for an observable p and a nonnegative integer d , is upward-closed with respect to \preceq_5^S . Furthermore, if \cong_5^S has finite index, then \preceq_5^S is a well-quasi-order. This is because $s \cong_5^S t$ implies $s \preceq_5^S t$: if there were an infinite sequence s_0, s_1, s_2, \dots of states such that for all $i \geq 0$ and $j < i$, we have $s_j \not\preceq_5^S s_i$, then no two of these states would be \cong_5^S equivalent.

(3 \Rightarrow 2) This part of the proof follows immediately from the stabilization property of well-quasi-orders [AČJT96]. \square

5.3 Decidable Properties: Bounded Reachability

Definition: Bounded-reachability logic The *bounded-reachability logic* L_5^\Diamond consists of the L_4^\Diamond -formulas that are generated by the grammar

$$\varphi ::= p \mid \varphi \vee \varphi \mid \exists \Diamond_{\leq d} \varphi \mid \exists \Diamond \varphi,$$

for constants $p \in \Pi$ and nonnegative integers d . The state logic \overline{L}_5^\Diamond consists of the duals of all L_5^\Diamond -formulas. \square

The following facts about bounded-reachability logic and its dual are relevant in our context. Both L_5^\Diamond and \overline{L}_5^\Diamond admit abstraction, and the state equivalence induced by both L_5^\Diamond and \overline{L}_5^\Diamond is \cong_5 (bounded-reach equivalence). It follows that the conjunction-free temporal logic L_4^\Diamond is more expressive than L_5^\Diamond , and \overline{L}_4^\Diamond is more expressive than \overline{L}_5^\Diamond . For example, the property that an observable can be reached in exactly d transitions can be expressed in L_4^\Diamond but not in L_5^\Diamond . Since L_5^\Diamond admits abstraction, and for STS5 systems the induced quotient can be constructed using the symbolic semi-algorithm **Reach**, we have the following theorem.

Theorem 5B *The L_5^\Diamond and \overline{L}_5^\Diamond model-checking problems are decidable for the class STS5 of symbolic transition systems.*

A direct symbolic model-checking semi-algorithm for L_5^\Diamond and, indeed, L_4^\Diamond is easily derived from the semi-algorithm **Reach**. Then, if **Reach** terminates, so does model checking for all L_4^\Diamond -formulas, including unbounded $\exists \Diamond$ properties. The extension to L_4^\Diamond is possible, because $\exists \bigcirc$ properties pose no threat to termination.

5.4 Example: Networks of Rectangular Hybrid Automata

A *network of timed automata* [AJ98] consists of a finite state controller and an arbitrarily large set of identical 1D timed automata. The continuous evolution of the system increases the values of all variables. The discrete transitions of the system are specified by a set of synchronization rules. We generalize the definition to rectangular automata. Formally, a *network of rectangular automata* is a triple (C, H, R) , where C is a finite set of controller locations, H is a 1D rectangular automaton, and R is a finite set of rules of the form $r = (\langle c, c' \rangle, e_1, \dots, e_n)$, where $c, c' \in C$ and e_1, \dots, e_n are jumps of H . The rule r is enabled if the




Fig. 9. Reach equivalence is coarser than bounded-reach equivalence

controller state is c and there are n rectangular automata H_1, \dots, H_n whose states are such that the jumps e_1, \dots, e_n , respectively, can be performed. The rule r is executed by simultaneously changing the controller state to c' and the state of each H_i , for $1 \leq i \leq n$, according to the jump e_i . The following result is proved in [AJ98] for networks of timed automata. The proof can be extended to rectangular automata using the observation that every rectangular automaton is simulated by an appropriate timed automaton [HKPV98].

Theorem 5C *The networks of rectangular automata belong to the class STS5. There is a network of timed automata that does not belong to STS4.*

6 General Symbolic Transition Systems

For studying reachability questions on symbolic transition systems, it is natural to consider the following fragment of bounded-reachability logic.

Definition: Reachability logic The *reachability logic* L_6^\diamond consists of the L_5^\diamond -formulas that are generated by the grammar

$$\varphi ::= p \mid \varphi \vee \varphi \mid \exists \Diamond \varphi,$$

for constants $p \in \Pi$. □

The reachability logic L_6^\diamond is less expressive than the bounded-reachability logic L_5^\diamond , because it induces the following state equivalence, \cong_6 , which is coarser than bounded-reach equivalence (see Figure 9: all states s_i , for $i \geq 0$, are reach equivalent, but no two of them are bounded-reach-equivalent).

Definition: Reach equivalence Let \mathcal{S} be a transition system. Two states s and t of \mathcal{S} are *reach equivalent*, denoted $s \cong_6^{\mathcal{S}} t$, if for every source- s trace of \mathcal{S} with target p , there is a source- t trace of \mathcal{S} with target p , and vice versa. The state equivalence \cong_6 is called *reach equivalence*. □

For every symbolic transition system \mathcal{R} with k observables, the reach-equivalence relation $\cong_6^{\mathcal{R}}$ has at most 2^k equivalence classes and, therefore, finite index. Since the reachability problem is undecidable for many kinds of symbolic transition systems (including Turing machines and polyhedral hybrid automata [ACH⁺95]), it follows that there cannot be a general algorithm for computing the reach-equivalence quotient of symbolic transition systems.

References

- [ACH⁺95] R. Alur, C. Courcoubetis, N. Halbwachs, T.A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995.
- [AČJT96] P. A. Abdulla, K. Čerāns, B. Jonsson, and Y.-K. Tsay. General decidability theorems for infinite-state systems. In *Proceedings of the 11th Annual Symposium on Logic in Computer Science*, pages 313–321. IEEE Computer Society Press, 1996.
- [AD94] R. Alur and D.L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.
- [AH98] R. Alur and T.A. Henzinger. *Computer-aided Verification: An Introduction to Model Building and Model Checking for Concurrent Systems*. Draft, 1998.
- [AHH96] R. Alur, T.A. Henzinger, and P.-H. Ho. Automatic symbolic verification of embedded systems. *IEEE Transactions on Software Engineering*, 22:181–201, 1996.
- [AJ98] P. Abdulla and B. Jonsson. Verifying networks of timed automata. In *TACAS 98: Tools and Algorithms for Construction and Analysis of Systems*, Lecture Notes in Computer Science 1384, pages 298–312. Springer-Verlag, 1998.
- [BFH90] A. Bouajjani, J.-C. Fernandez, and N. Halbwachs. Minimal model generation. In *CAV 90: Computer-aided Verification*, Lecture Notes in Computer Science 531, pages 197–203. Springer-Verlag, 1990.
- [Dam94] M. Dam. CTL^{*} and ECTL^{*} as fragments of the modal μ -calculus. *Theoretical Computer Science*, 126:77–96, 1994.
- [EJS93] E.A. Emerson, C.S. Jutla, and A.P. Sistla. On model checking for fragments of μ -calculus. In *CAV 93: Computer-aided Verification*, Lecture Notes in Computer Science 697, pages 385–396. Springer-Verlag, 1993.
- [FS98] A. Finkel and Ph. Schnoebelen. *Well-structured Transition Systems Everywhere*. Technical Report LSV-98-4, Laboratoire Spécification et Vérification, ENS Cachan, 1998.
- [Hen95] T.A. Henzinger. Hybrid automata with finite bisimulations. In *ICALP 95: Automata, Languages, and Programming*, Lecture Notes in Computer Science 944, pages 324–335. Springer-Verlag, 1995.
- [Hen96] T.A. Henzinger. The theory of hybrid automata. In *Proceedings of the 11th Annual Symposium on Logic in Computer Science*, pages 278–292. IEEE Computer Society Press, 1996.
- [HHK95] M.R. Henzinger, T.A. Henzinger, and P.W. Kopke. Computing simulations on finite and infinite graphs. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 453–462. IEEE Computer Society Press, 1995.
- [HHWT95] T.A. Henzinger, P.-H. Ho, and H. Wong-Toi. HYTECH: the next generation. In *Proceedings of the 16th Annual Real-time Systems Symposium*, pages 56–65. IEEE Computer Society Press, 1995.
- [HK96] T.A. Henzinger and P.W. Kopke. State equivalences for rectangular hybrid automata. In *CONCUR 96: Concurrency Theory*, Lecture Notes in Computer Science 1119, pages 530–545. Springer-Verlag, 1996.
- [HKPV98] T.A. Henzinger, P.W. Kopke, A. Puri, and P. Varaiya. What’s decidable about hybrid automata? *Journal of Computer and System Sciences*, 57:94–124, 1998.
- [HM99] T.A. Henzinger and R. Majumdar. Symbolic model checking for rectangular hybrid systems. Submitted for publication, 1999.
- [KS90] P.C. Kanellakis and S.A. Smolka. CCS expressions, finite-state processes, and three problems of equivalence. *Information and Computation*, 86:43–68, 1990.
- [vG90] R.J. van Glabbeek. *Comparative Concurrency Semantics and Refinement of Actions*. PhD thesis, Vrije Universiteit te Amsterdam, The Netherlands, 1990.